

**APPLICATION FOR
UNITED STATES LETTERS PATENT
FOR**

METHOD AND APPARATUS FOR SHARED FLOW CONTROL OF DATA

By

**Larry Lee Clark, Michael Lawrence Denburg,
Stephen Ray Moore and Jeffri Harold Frontz**

CERTIFICATE OF EXPRESS MAIL

"Express Mail" mailing label number: EE 434581932US
Date of Deposit: MAY 24, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" on the date indicated above and is addressed to the Assistant Commissioner For Patents, Washington, D.C. 20231.

WEAVERA PETROWSKI



METHOD AND APPARATUS FOR SHARED FLOW CONTROL OF DATA

BACKGROUND OF THE INVENTION

This invention relates to communication of data, and in particular, to the upstream and downstream flow control of data.

5 Layered communication protocols are often implemented as a series of sequential functional layers. A well known example of such a sequential process is the seven layer Open System Interconnection (OSI) protocol model. Each software layer in the OSI model has a protocol provider that implements a portion
10 of the overall protocol stack. A software protocol provider typically processes the data passing through its layer as fast as possible.

Traditional flow control approaches require each end of the communication link to implement a flow control mechanism. For
15 example, the X.25 and TCP protocol providers use peer-to-peer flow control to govern the data exchange. However, the high bandwidth of the underlying transmission medium can allow a very large amount of data to be exchanged even though it is being regulated by the existing protocol layers. This high data rate
20 can exhaust or deplete a machine's resources to the point that the machine may not be able to perform other more important tasks.

A method and apparatus to enforce additional data flow

control is needed in a single flow control layer that is independently accessible (shared) by the transmitting and the receiving devices or processes.

5

SUMMARY OF THE INVENTION

A flow control module is placed in the streams of data as it passes between a plurality of sources from above and below the flow control module. The flow control module calculates via aggregate counters the aggregate data rate for the upstream and downstream streams of data and the individual data rate (bytes per second) via individual data rate counters associated with each stream of data in both the upstream stream direction and the downstream directions of the streams of data. The aggregate counters are compared to predetermined threshold values for the upstream and downstream data and throttled appropriately. Additionally, the individual data rate counters are each compared to a predetermined individual data rate threshold value and throttled appropriately. The flow control of data using the aggregate data rate counters and the individual data rate counters act upon the data in the upstream and downstream when the threshold values are exceeded. Furthermore, the counters are decremented at predetermined periods by a predetermined number of bytes.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing objects and advantageous features of the invention will be explained in greater detail and others will be made apparent from the detailed description of the present invention, which is given with reference to the several figures of the drawing, in which:

Fig. 1 illustrates a block diagram of a flow control module coupled to a plurality of applications and a protocol provider in accordance with an embodiment of the invention;

Fig. 2 is a flow diagram illustration of a method of shared flow control of streams of data in the downstream from applications to a network in accordance with an embodiment of the invention;

Fig. 3 is a flow diagram illustration of a method enabling streams of data in the downstream direction that have been disabled in accordance with an embodiment of the invention;

Fig. 4 is a flow diagram illustration of a method of shared flow control of streams of data in the upstream from the transport layer provider to applications in accordance with an embodiment of the invention;

Fig. 5 is a flow diagram illustration of a method enabling streams of data in the upstream direction that have been disabled in accordance with an embodiment of the invention; and

Fig. 6 is a system block diagram of a application process

and a network communicating in accordance with an embodiment of the invention.

DETAILED DESCRIPTION

5 In Fig. 1, an illustration of a block diagram of a flow control module 102 coupled to a plurality of applications 104-110 and a transport layer provider 112 in accordance with an embodiment of the invention. The flow control module 102 is coupled through software to a plurality of applications, labeled
10 "APPLICATION-1" to "APPLICATION-N", 104-110 respectively. The applications 104-110 are contained at the application software level and interface with the operating system software level. The flow control module 102 is also coupled to the transport layer provider 112 within the operating system software level.
15 The transport layer provider 112 is coupled via software and device drivers to hardware that is connected to a physical media 114 associated with a network.

An aggregate stream of data is comprised of a plurality of individual streams of data from applications 104-110. The flow
20 control module 102 adds to or increments an aggregate counter by the amount of data (i.e. bytes of data) that is received at the flow control module 102 per period of time (i.e. second). The aggregate counter is checked every time data is encountered by comparing the aggregate counter value to a predetermined

aggregate threshold (80,000 byte). The aggregate counter is decremented by a predetermined amount every time a timer expires. The timer is set to a predetermined value, such as 1 second. If the aggregate threshold is exceeded when the
5 aggregate counter value is checked then all streams of data from the applications 104-110 are throttled.

Similarly, the individual data rate of each stream of data that make up the combined aggregate stream of data is determined using an individual counter associated with the individual
10 stream of data. If the individual counter exceeds a predetermined threshold value of data bytes, then the stream of data associated with the individual counter is throttled. The individual counters are decremented in a similar way as the aggregate counters. At the expiration of a timer having a
15 predetermined period, such as one second, each of the individual counter values are decremented by a predetermined amount (8,000 bytes). The amount to decrement each individual counter can be unique for each individual counter or all individual counters can be decremented by the same value (8,000 bytes).

20 Furthermore, the individual and aggregate counters in the present embodiment can never go below zero. In an embodiment of the invention the UNIX operating system using the UNIX streams library may selectively be used to enable and disable data flow in response to the individual and aggregate predetermined

thresholds.

The applications 104-110 are also active and may be receiving or transmitting data at the same time. As described above, the flow control module 102 has an upstream and downstream aggregate counter for counting the received and transmitted bytes from applications 104-110. When the upstream or downstream aggregate counter exceeds the associated aggregate data rate threshold for a period of time (i.e. a second), the flow control module throttles all the streams of data from the applications 104-110. Similarly, when the individual stream of data exceeds the individual data rate threshold, then the individual stream of data is throttled.

When the aggregate data rate threshold is no longer exceeded by the aggregate counter, then the throttling of the aggregate stream of data ceases, provided the individual data rate threshold is not exceeded. Similarly, the individual stream of data is unthrottled once the individual data rate for the individual stream of data no longer exceeds the individual data rate threshold.

Turning to Fig. 2, a flow diagram illustration of a method of shared flow control of streams of data in the downstream from applications 104-110, Fig. 1, to a network in accordance with an embodiment of the invention. In step 202, Fig. 2, a plurality of streams of data are received at the flow control module 102,

Fig. 1, from the applications 104-110. A downstream aggregate counter in the flow control module 102 is incremented by the amount of data (number of bytes) received in step 204, Fig. 2. Similarly, in step 206, the individual downstream counters
 5 associated with the individual downstream streams of data are incremented. In step 208, the data received in step 202 is transmitted from the flow control module 102, Fig. 1.

The downstream aggregate threshold is compared to the aggregate downstream counter in step 210, Fig. 2. If the
 10 downstream aggregate threshold has been exceeded in step 210, then in step 212 all the streams of data in the downstream direction are throttled and processing is halted in 218.

If in step 210, the downstream aggregate threshold has not been exceeded then the individual downstream counters are
 15 compared to associated individual downstream thresholds in step 214. If no individual data rate threshold is exceeded in step 214, then step 202 is repeated. If one or more of the individual data rate threshold are exceeded in step 214, then in step 216 the data streams that exceeded the individual threshold
 20 are throttled and then Step 202 is repeated.

In Fig. 3, a flow diagram illustration of a method enabling streams of data in the downstream direction that have been disabled in accordance with an embodiment of the invention is shown. In step 302, the downstream aggregate counter is

decremented by a predetermined amount (80,000 bytes). The individual downstream counters are decremented by a predetermined amount (8,000 bytes) in step 304. In step 306, the aggregate downstream counter is compared to the downstream aggregate threshold. If the aggregate downstream counter is less than the downstream aggregate threshold in step 306, then each of the individual downstream data streams that are less than their associated individual thresholds are unthrottled in step 308 and further processing waits for the decrement timer to expire in step 310. Otherwise, the aggregate downstream counter equals or exceeds the downstream aggregate threshold in step 306 and in step 310 processing waits for the decrement timer to expire. The decrement timer is set to a predetermined value, such as one second, upon initialization.

In Fig. 4, a flow diagram illustration of a method of shared flow control of streams of data in the upstream from the transport layer provider 112, Fig. 1, to applications 104-110 in accordance with an embodiment of the invention is shown. In step 402, Fig. 4, a plurality of streams of data are received at the flow control module 102, Fig. 1, from the transport layer provider 112. An upstream aggregate counter in the flow control module 102 is incremented by the amount of data (number of bytes) received in step 404, Fig. 4. Similarly, in step 406, the individual upstream counters associated with each of the

individual upstream streams of data are incremented. In step 408, the data received in step 402 is transmitted from the flow control module 102, Fig. 1.

The upstream aggregate threshold (80,000 bytes) is compared
5 the aggregate upstream counter in step 410, Fig. 4. If the upstream aggregate threshold has been exceeded in step 410, then in step 412 all the streams of data in the upstream are throttled and processing is halted in 418.

If in step 410, the upstream aggregate threshold has not
10 been exceeded, then the individual upstream counters are compared to associated individual upstream thresholds in step 414. If no individual data rate threshold is exceeded in step 414, then step 402 is repeated. If one or more of the individual data rate threshold are exceeded in step 414, then in
15 step 416 the streams of data that did exceed the threshold are throttled and step 402 is repeated.

In Fig. 5, a flow diagram illustration of a method enabling streams of data in the upstream direction that have been disabled in accordance with an embodiment of the invention is
20 shown. In step 502, the upstream aggregate counter is decremented by a predetermined amount (80,000 bytes). Similarly, the individual upstream counters are decremented by a predetermined amount (8,000 bytes) in step 504. In step 506, the aggregate upstream counter is compared to the upstream

aggregate threshold. If the aggregate upstream counter is less than the upstream aggregate threshold in step 506, then each of the individual upstream streams of data that are less than their associated individual data rate are unthrottled in step 508 and further processing waits for the decrement timer to expire in step 510. Otherwise, the aggregate upstream counter equals or exceeds the upstream aggregate threshold in step 506, then in step 510, processing waits for the decrement timer to expire. The decrement timer is set to a predetermined value, such as one second, upon initialization.

The steps illustrated in Fig. 2 through Fig. 5 may selectively be implemented by computer-readable program code (i.e. UNIX with "C" or "C++") or in integrated circuits. The computer-readable program code is contained in an article of manufacture such as a compact disk read only memory (CD ROM), floppy disk, magnetic tape, or programmable chip (i.e. PROM, EEPROM, or EPROM).

In Fig. 6, a system block diagram of a application process and a network communicating is shown. The communication flows to and from the network. A plurality of data streams 602 between applications 604 and a network 606 pass (in both the upstream and downstream) through a flow control module 608 having a transmitter/receiver 610, a counter 612, and a comparator 614. The data stream 602 also passes through a

transport layer provider 416 having an implementor 618 and another transmitter/transmitter 620.

A data stream 402 between the application processes 604 and the network 606 have a data rate that is calculated in the flow control module 608. The flow control module 608 receives the data stream at the transmitter/receiver 610 and counters 612 having an upstream aggregate counter and downstream aggregate counter determine the data rate for the aggregate stream of data. The upstream and downstream aggregate counters are compared to an aggregate data rate threshold by the comparator 614. In an alternate embodiment, multiple data streams each have an associated data rate threshold that is calculated by an individual upstream counter and individual downstream counter. The individual counters are compared by the comparator 614 to an individual data rate threshold.

If either of the thresholds (aggregate or individual) is exceeded, then the data rate of a stream of data is throttled. The flow control module 608 signals to the implementor 618 via a signal (electrical or software) to throttle the data dream. The implementor 618 in the transport layer provider 616 then throttles the data rate of the stream of data that is being sent to the network 606 by the other transmitter/receiver 620. Therefore, the machine resources are conserved by managing the data flow per unit of time with a shared flow control module

While the invention has been particularly shown and described with reference to a particular embodiment, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the spirit and scope of the invention and it is intended that all such changes come within the scope of the following claims.

[illegible]